

Cross-dataset Deep Transfer Learning for Activity Recognition

Martin Gjoreski

Jožef Stefan Institute, Department of Intelligent Systems,
Jožef Stefan Postgraduate School, Slovenia,
martin.gjoreski@ijs.si

Mitja Luštrek

Matjaž Gams

Jožef Stefan Institute, Department of Intelligent Systems,
Jožef Stefan Postgraduate School, Slovenia,
martin.gjoreski@ijs.si

Stefan Kalabakov

Faculty of Electrical Engineering and Information
Technologies,
University Ss. Cyril and Methodius, N. Macedonia

Hristijan Gjoreski

Faculty of Electrical Engineering and Information
Technologies,
University Ss. Cyril and Methodius, N. Macedonia

ABSTRACT

Convolution Neural Network (CNN) filters learned on one domain can be used as feature extractors on another similar domain. Transferring filters allow reusing datasets across domains and reducing labelling costs. In this paper, four activity recognition datasets were analyzed to study the effects of transferring filters across the datasets. A spectro-temporal ResNet was implemented as a deep, end-to-end learning architecture. We analyzed the number of transferred CNN residual blocks with respect to the size of the target-adaptation data. The analysis showed that transfer learning using small adaptation subsets is more useful when the target domain contains a small number of different activities. Furthermore, the similarity between the domains, participating in the transfer learning scenario, seems to play a role in its success. The most successful transfer achieved an F1-score of 93%, which is an increase of 9 percentage points compared to a domain-specific baseline model.

CCS CONCEPTS

• Human-centered computing • Human computer interaction (HCI) • Ubiquitous and mobile computing

KEYWORDS

Activity recognition, Deep learning, End-to-end learning, Transfer learning

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

UbiComp/ISWC '19 Adjunct, September 9–13, 2019, London, United Kingdom
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6869-8/19/09...\$15.00
<https://doi.org/10.1145/3341162.3344865>

ACM Reference format:

Martin Gjoreski, Stefan Kalabakov, Mitja Luštrek, Majaz Gams and Hristijan Gjoreski. 2019. Cross-dataset Deep Transfer Learning for Activity Recognition. In *Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and the 2019 International Symposium on Wearable Computers (UbiComp/ISWC '19 Adjunct)*, September 9–13, 2019, London, United Kingdom. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3341162.3344865>

1 INTRODUCTION

Deep learning represents a class of machine learning algorithms that use a cascade of multiple layers of non-linear processing units [1]. The first layer receives the input data, and each successive layer uses the output from the previous layer as the input. The basic idea dates from 1943 when McCulloch and Pitts created the first computational model of neural networks (NN) based on threshold logic [2]. Fast forward to the modern era, where large processing power and memory storage are relatively affordable, deep learning architectures are used to solve complicated AI tasks (e.g., in computer vision, language, biomedicine, etc.) by learning high-level abstractions [3][4].

In the last decade, many attempts at activity recognition (AR) were made with DL end-to-end architectures [8][9]. The focus was on Convolutional Neural Networks (CNNs). The CNNs are a type of NNs that are specially designed to ensure some degree of shift, scale and distortion invariance [5][6][7]. In addition, the CNNs can automatically capture hierarchical feature representations of the data [7]. Thus, CNNs, in combination with subsampling layers (e.g., pooling layers) and fully connected layers, are a very powerful end-to-end learning architecture [10][11][12]. However, because of the specific architecture (parameter sharing and local connections), the CNNs have much fewer connections and parameters to train compared to other types of NNs, e.g., fully connected NNs.

Since classical (feature-based) AR requires substantial amounts of labeled training data to perform well under diverse circumstances [16], end-to-end learning architectures require even bigger amounts of labeled data. One such example is the DeepConvLSTM introduced by Ordonez et al. [13] in 2016. The DeepConvLSTM uses several CNN layers stacked over a long short-term memory (LSTM) recurrent neural network (RNN) for AR [14]. To reduce the costs of acquiring labeled data, transfer learning can be employed. In the AR domain, transfer learning techniques provide mechanisms for transferring instances, feature-representations, model-parameters and/or transfer or relational-knowledge [17]. Although transfer learning has been demonstrated to be feasible [18][19][32], it still remains a challenging task. Evidence from other domains (e.g., computer vision), where CNNs have been successfully applied for supervised tasks, suggest that the filters learned by the CNNs on one domain can be used as feature extractors on another similar domain. By analogy, this form of transfer learning has been also analyzed in the AR domain in the recent years. Ding et al. [20] analyzed deep transfer learning between users with unlabeled data using an end-to-end domain-adversarial neural network based on CNNs. Their architecture used shared filters across all sensor modalities, i.e., acceleration, magnitude and gyroscope data, applied over the temporal representation of the signals. Morales et al. [21] analyzed the DeepConvLSTM for transfer learning across mobile AR datasets, sensor modalities and sensor locations. Among other things, the authors have analyzed whether the features learned by the CNN’s filters on the OPPORTUNITY dataset [22] can be re-used for building a model for the Skoda [23]. Their study confirmed that transferring filters across datasets is quite challenging and in most of the cases the domain-specific baseline model outperformed the models build using transfer learning. Similarly as Ding et al. [20], the DeepConvLSTM uses shared filters across all sensor channels, i.e., acceleration data, applied over the temporal representation of the signals.

In this study, we analyzed a deep multimodal spectro-temporal ResNet (Multi-ResNet), which has already been proved successful for AR in our previous study [27] by achieving comparable accuracy to state-of-the art feature-based models on one of the largest AR datasets, the SHL dataset [29]. The Multi-ResNet is explained in Section 3. Compared to the related work, our study provides several novelties: (i) to the best of our knowledge, this is the first study that analyzes four different datasets for transfer learning in the AR domain; (ii) our DNN architecture provides the possibility to learn CNN filters individually for each channel and thus to transfer channel-specific CNN filters; (iii) our DNN architecture provides the possibility to learn CNN filters for both temporal and spectral representations of raw signals, and thus it provides the possibility to transfer CNN filters not only in the temporal but also in the spectral domain.

2 DATASETS

The four datasets used to evaluate the effectiveness of the transfer learning method are: Skoda [23] OPPORTUNITY [22], JSI-FOS [24][25] and PAMAP2 [26]. All datasets, except for Skoda, are comprised of activities of daily living, whereas Skoda is comprised of activities performed by a worker at a quality control checkpoint in a car factory. All of them contain information captured by a 3D accelerometer on the dominant hand of the wearer, which makes it ideal to use in this analysis since it focuses our transfer learning methods to the domain of the task, rather than the modality or location of the sensors used. More detailed information on the datasets can be found in Table 1.

Table 1: Overview of the experimental datasets.

Dataset	Type	#Subjects	Sampling Rate	#Activities	Duration [h]
OPPORTUNITY [22]	ADL	4	~30Hz	21	10.6
Skoda Mini Checkpoint [23]	Factory	1	96Hz	10	2.9
PAMAP2 [26]	ADL	9	100Hz	15	15.1
JSI-FOS [24][25]	ADL	10	50Hz	8	40.2

3 DEEP LEARNING ARCHITECTURE

The DNN architecture (Figure 1) used in this study is a deep multimodal spectro-temporal ResNet (Multi-ResNet). The Multi-ResNet has already been proved successful for AR in our previous study [27] by achieving comparable accuracy to state-of-the art feature-based models on one of the largest AR datasets, the SHL dataset [29]. The structure is based on an idea for training very deep end-to-end networks for image recognition: it uses shortcut (residual) connections to fight the gradient-vanishing problem [28]. Additionally, the network has two key factors for a successful AR system – it utilizes multimodal and spectro-temporal information. For each sensor channel, the network extracts channel-specific spectro-temporal information: the spectral information is extracted by calculating a spectrogram of the input signal, which is then used as input to a 1D convolutional layer; the temporal representation is extracted by residual blocks that contain 1D CNN filters. To reduce internal covariance shift, each CNN layer is followed by a batch normalization (BN) layer [30]. To speed up the training process, ReLU activation layers are used [31]. For dimensionality reduction, each residual block ends up with a maximum pooling layer. The output of the channel-specific layers, i.e., channel-specific spectro-temporal information, is then merged by two dense (fully connected) layers. To avoid overfitting, L2 regularization and dropout were used for the dense layers. The final output of the Multi-ResNet is provided by a softmax layer, which outputs probabilities for each class.

4 EXPERIMENTS

Experimental Setup

Four datasets were used in this study and the effects of transfer learning were tested on each combination of these datasets. Firstly, all datasets were downsampled to a frequency of 25Hz and segmented to windows of 4s with an overlap between windows of

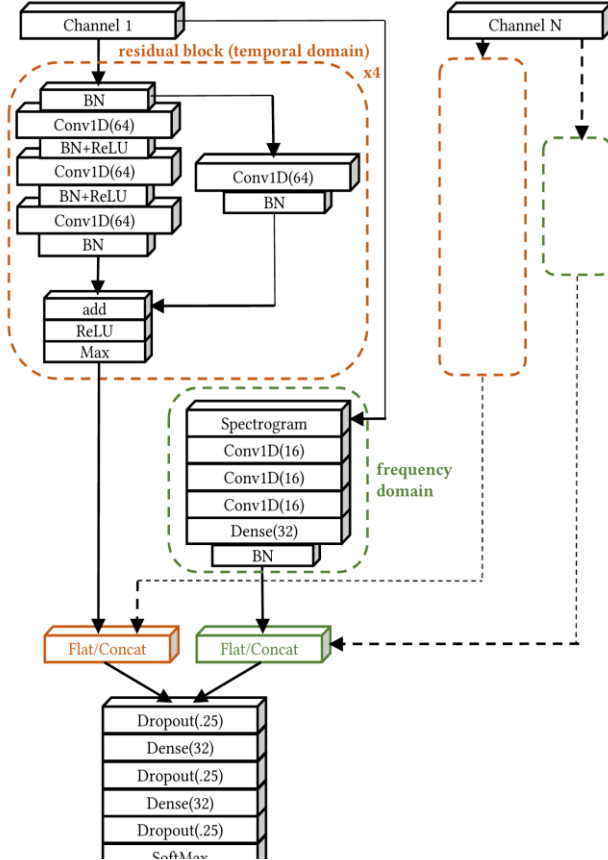


Figure 1: The deep multimodal spectro-temporal ResNet (Multi-ResNet) used in the study.

2s. Furthermore, all datasets were split into a train, validation and test subsets, which comprised 40%, 10% and 50% of the overall duration of each dataset, respectively.

To get more realistic evaluation of our work, the split was done in such a way that every subject would appear in only one of the subsets. In the case where there was only one subject in the dataset, such as in the Skoda dataset, the split was done per activity occurrence so that neighboring instances could occur in only one of the subsets.

In each experiment, one of the datasets served the purpose of a source domain and another dataset served as a target domain. At the start of each experiment, a model (source model) was trained on the training subset of the source dataset and validated on the validation subset of that same dataset. The training was

fully supervised, by back-propagating the gradients through all the layers. In the next step, the model weights were transferred from the source model to a new model (target model), which would then go on to be adapted (trained) on a portion of the training subset of the target dataset (adaptation subset). In this case, all weights of the source model were transferred to the target model and, depending on the experiment, some of them were frozen during the training of the network. Thus, the gradients were back-propagated only up to a certain layer of the target model.

The experiments were carried out in such a way that the weights of whole residual blocks were transferred from the source model to the target model and not just individual convolutional layers. Also, several different sizes of the adaptation subset were tested in these experiments. However, the maximum number of instances never exceeded the full size of the training subset of the Skoda dataset, which is the smallest dataset of all four. Thus, in summary, two parameters were quantified in the experiments:

- (i) the number of residual blocks whose weights would be transferred from the source model to the target model;
- (ii) the size of the adaptation subset on which the target model would adapt to the new domain.

In addition, each experiment included the results of a model which was trained solely on the adaptation subset and had no weights transferred to it from the source model. This domain-specific model served as a baseline to which we compared our results. Finally, the results were scored using the F1-score with micro average.

All models were trained by minimizing the cross-entropy loss function using the Adam optimizer, using a learning rate of 10^{-3} and a decay of 10^{-3} . The batch size was set to 256 and the maximum number of training epochs was set to 70. The network parameters, including the number of residual blocks, the number of CNN layers per block, the size of the CNN filters, the learning rate and the batch size, were determined experimentally.

Experimental Results

The results of our experiments can be seen in Figure 2. The title of each heatmap represents: which dataset was used as the target domain (T); in the brackets is the F1-score achieved by a majority classifier for the specific target domain (13 for Skoda, 19 for OPORTUNITY, 26 for PAMAP and 24 for JSI-FOS); which dataset was used as the source domain (S). The vertical axis of every heatmap represents the number of instances in the adaptation subset. The horizontal axis represents the number of residual blocks whose weights were transferred from the source model to the target model. The numbers in each cell represents the micro

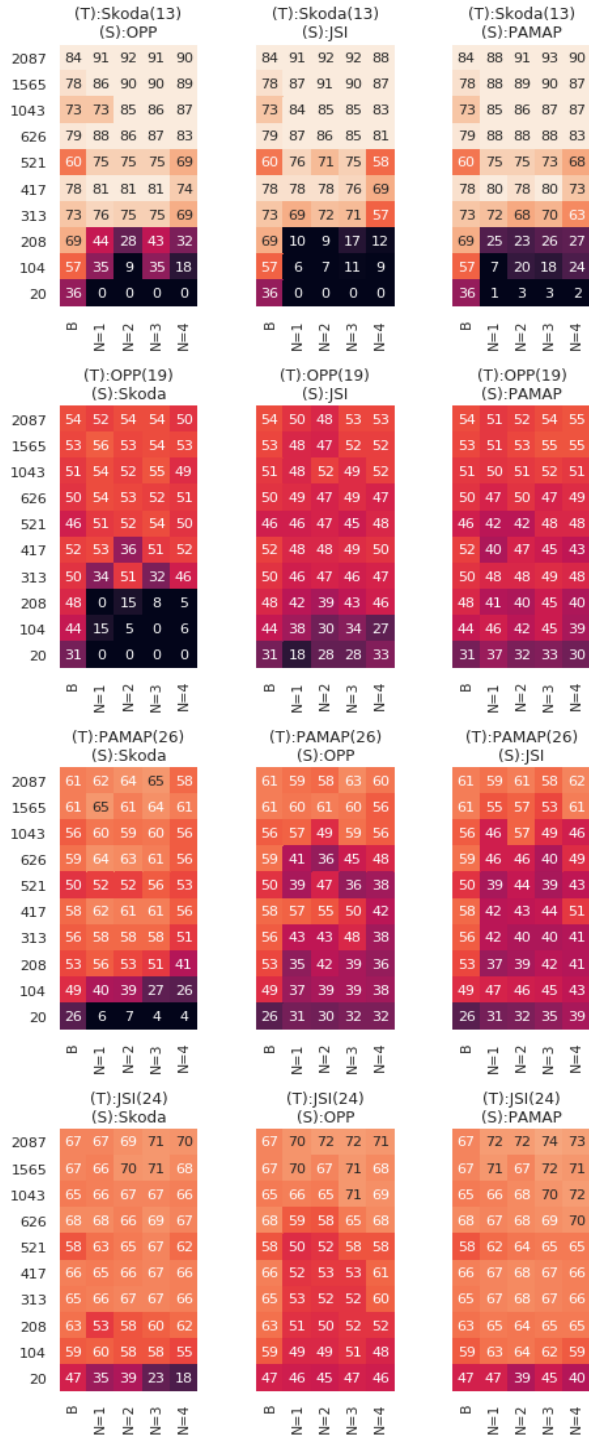


Figure 2: F1-score for each Target (T) - Source (S) domain combination. The y-axis represents the number of instances in the adaptation subset. The x-axis represents the number of residual blocks transferred from the S model to the T model. The leftmost column shows the F1-score of the domain-specific baseline model.

F1-score achieved for the specific combination of number of instances in the adaptation set and the number of residual blocks whose weights were transferred. The leftmost column shows the F1-score of the domain-specific baseline model which was trained on that same adaptation subset of the target domain.

When Skoda is the target dataset (first-row heatmaps in Figure 2) and the adaptation subset contains more than 400 instances, we can observe an increase of performance compared to the baseline in all transfer scenarios. When the number of instances in the adaptation subset is less than 400, there is a significant drop in the F1 scores, with OPPORTUNITY proving to be the best choice of a source dataset in the situation.

When OPPORTUNITY is the target dataset (second-row heatmaps in Figure 2), it can be noted that no matter which of the datasets is used as source, the transfer learning scenarios either bring a decrease of performance or in the best case are on par with the baseline score. Additionally, when Skoda appears as a source dataset and the adaptation subset is very small, a severe degradation of performance can be observed compared to the baseline and the scenarios in which PAMAP and JSI-FOS are used as source datasets.

When JSI-FOS is the target dataset (last-row heatmaps in Figure 2), another increase of performance can be seen when Skoda or PAMAP are used as the source datasets. More specifically, the transfer from PAMAP to JSI-FOS provides more consistent results. When OPPORTUNITY is the source dataset and the adaptation dataset is large enough, some increases of performance can also be observed.

Finally, when PAMAP is the target dataset (third-row heatmaps in Figure 2) and JSI-FOS or OPPORTUNITY are used as source datasets, the scores for the baseline and transfer learning, with any number of weights transferred, are within a few percentage points of each other when we use the maximum or near to the maximum number of instances in the adaptation subset. In all other cases these sources provide a decrease of performance. When Skoda is used as the source dataset, transfer learning provides an increase of performance, except in the case of an extremely small adaptation dataset, with less than 200 instances.

5 CONCLUSION AND DISCUSSION

In our study, we looked at four different domains as sources and targets of our transfer learning experiments. The results show that the number of different activities that occur in the target dataset is of great importance. This is especially true because the transfer learning scenarios show performance improvement only when JSI-FOS or Skoda appear as target datasets. JSI-FOS and Skoda include eight and ten different activities, respectively, while OPPORTUNITY and PAMAP include 21 and 15 different activities, respectively. This points to the conclusion that small adaptation

subsets do not provide enough data that the models could use to distinguish between a high number of different activities, such as in the case where OPPORTUNITY and PAMAP are used as target datasets.

In those scenarios where transfer learning produced an increase of performance, the similarity of the datasets involved seemed to also play a big role. OPPORTUNITY and Skoda as datasets are comprised of fine-grained gestures and contain a lot of diversity, as opposed to JSI-FOS and PAMAP, which are comprised of more general, high-level activities. The importance of this similarity can be observed especially in the case of an extremely small adaptation subset when transferring to Skoda as a target dataset, where using OPPORTUNITY as a source yields the most consistent results. Another example of this is the consistency in the results of the transfer learning models when transferring from PAMAP to JSI-FOS.

Our experiments focused on building end-to-end models with very little data, i.e., starting from 20 instances (near 1 minute of data) to 2087 instances (near 75 minutes of data). As a consequence, the overall accuracy varies between 48% (when transferring to OPPORTUNITY) to 93% when transferring to Skoda. Also, the experimental results are quite related to the specific DNN architecture, the Multi-ResNet. In future, we plan to test additional end-to-end DL architectures and to compare their performance. Also, we plan to use multiple datasets as a source, and this way to learn more general model. Finally, we plan to provide a comparison of transfer learning performed in the spectral domain as opposed to a transfer learning in the temporal domain.

REFERENCES

- [1] L. Deng, D. Yu. (2014). Deep Learning: Methods and Applications (PDF). Foundations and Trends in Signal Processing. 7 (3–4): 1–199. doi:10.1561/20000000039.
- [2] W. McCulloch, W. Pitts (1943). A Logical Calculus of Ideas Immanent in Nervous Activity. Bulletin of Mathematical Biophysics. 5 (4): 115–133. doi:10.1007/BF02478259.
- [3] Y. Bengio (2009). "Learning Deep Architectures for AI." Technical Report 1312. [Online]. Accessible at: <https://www.iro.umontreal.ca/~lisa/pointeurs/TR1312.pdf>
- [4] A. Krizhevsky, I. Sutskever, G. E. Hinton (2012). ImageNet classification with deep convolutional neural networks. in International Conference on Neural Information Processing Systems, pp. 1097–1105.
- [5] H. Lee, R. Grosse, R. Ranganath, A.Y. Ng (2009). Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In Proceedings of the 26th Annual International Conference on Machine Learning, pages 609–616. ACM.
- [6] K. Jarrett, K. Kavukcuoglu, M. A. Ranzato, and Y. LeCun (2009). What is the best multi-stage architecture for object recognition? In International Conference on Computer Vision, pages 2146–2153. IEEE.
- [7] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner (1998). Gradient-based learning applied to document recognition. Proceedings of the IEEE, vol. 86, no. 11, pp. 2278–2324.
- [8] H. Gjoreski, J. Bizjak, M. Gjoreski, M. Gams (2016). Comparing deep and classical machine learning methods for human activity recognition using wrist accelerometer, in: Proceedings of the IJCAI 2016 Workshop on Deep Learning for Artificial Intelligence, New York, NY, USA, Vol. 10.
- [9] D. Ravi, C. Wong, B. Lo, G.-Z. Yang (2017). A deep learning approach to on-node sensor data analytics for mobile or wearable devices, IEEE journal of biomedical and health informatics 21 (1) 56–64.
- [10] S. Bhattacharya, N. D. Lane, Sparsification and separation of deep learning layers for constrained resource inference on wearables, in: Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM, ACM, 2016, pp. 176–189.
- [11] M. Zeng, L. T. Nguyen, B. Yu, O. J. Mengshoel, J. Zhu, P. Wu, J. Zhang (2014). Convolutional neural networks for human activity recognition using mobile sensors, in: 6th International Conference on Mobile Computing, Applications and Services, IEEE, pp. 197–205.
- [12] J. Yang, M. N. Nguyen, P. P. San, X. L. Li, S. Krishnaswamy (2015). Deep convolutional neural networks on multichannel time series for human activity recognition, in: Twenty-Fourth International Joint Conference on Artificial Intelligence.
- [13] F. Ordonez, D. Roggen (2016). Deep convolutional and LSTM recurrent neural networks for multimodal wearable activity recognition, Sensors 16 (1) 115.
- [14] S. Hochreiter, J. Schmidhuber (1997). Long short-term memory, Neural computation 9 (8) 1735–1780.
- [15] K. Krishna, D. Jain, S. V. Mehta, S. Choudhary (2018). An LSTM based system for prediction of human activities with durations, Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies 1 (4) 147.
- [16] D. Cook, K. D. Feuz, N.C. Krishnan (2013). Transfer learning for activity recognition: A survey. Knowledge and information systems, 36(3), pp.537–556.
- [17] S. Pan, Q. Yang (2010). A survey on transfer learning. Knowledge Data Engineering IEEE Trans 22(10):1345–1359
- [18] A. Calatroni, D. Roggen, and G. Tröster (2011). Automatic transfer of activity recognition capabilities between body-worn motion sensors: Training newcomers to recognize locomotion. In Proc. INSS, Vol. 6.
- [19] M. Kurz, M. Hölzle, A. Ferscha, A. Calatroni, and others (2011). Real-time transfer and evaluation of activity recognition capabilities in an opportunistic system. In Proc. ADAPTIVE. 73–78.
- [20] R. Ding et al. (2019). Empirical Study and Improvement on Deep Transfer Learning for Human Activity Recognition. Sensors, 19(1), 57.
- [21] F. J. O. Morales, D. Roggen, (2016). Deep convolutional feature transfer across mobile activity recognition domains, sensor modalities and locations. In Proceedings of the 2016 ACM International Symposium on Wearable Computers (pp. 92–99). ACM.
- [22] D. Roggen, et al. (2010). Collecting complex activity datasets in highly rich networked sensor environments. In 2010 Seventh international conference on networked sensing systems (INSS) (pp. 233–240). IEEE.
- [23] P. Zappi et al. (2008). Activity recognition from on-body sensors: accuracy-power trade-off by dynamic sensor selection. In European Conference on Wireless Sensor Networks (pp. 17–33). Springer, Berlin, Heidelberg.
- [24] H. Gjoreski, B. Kaluža, M. Gams, Radoje Milić, M. Luštrek (2015). Context-based Ensemble Method for Human Energy Expenditure Estimation. Applied Soft Computing, Pages 960–970.
- [25] S. Kozina, H. Gjoreski, M. Gams, M. Luštrek (2013). Three-layer activity recognition combining domain knowledge and meta-classification. Journal of Medical and Biological Engineering, vol. 33, no. 4. p.p. 406–414.
- [26] A. Reiss, D. Stricker (2012). Introducing a new benchmarked dataset for activity monitoring. In 2012 16th International Symposium on Wearable Computers (pp. 108–109). IEEE.
- [27] M. Gjoreski et al. (2019). Classical and deep learning methods for recognizing human activities and modes of transportation with smartphone sensors. Information Fusion, Under review.
- [28] S. Hochreiter, Y. Bengio, P. Frasconi, J. Schmidhuber, et al. (2001). Gradient flow in recurrent nets: the difficulty of learning long-term dependencies
- [29] H. Gjoreski, M. Ciliberto, L. Wang, F. J. O. Morales, S. Mekki, S. Valentin, D. Roggen (2018). The University of Sussex-Huawei Locomotion and Transportation Dataset for Multimodal Analytics with Mobile Devices." IEEE Access, 2018, DOI: 10.1109/ACCESS.2018.2858933
- [30] S. Ioffe, C. Szegedy (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift, arXiv preprint arXiv:1502.03167.
- [31] V. Nair, G. E. Hinton (2010). Rectified linear units improve restricted Boltzmann machines, in: Proceedings of the 27th international conference on machine learning (ICML-10), 2010, pp. 807–814.
- [32] S. Inoue, X Pan. (2016). Supervised and Unsupervised Transfer Learning for Activity Recognition from Simple In-home Sensors. MobiQuitous.